Exhibit A

NUANCE COMMUNICATIONS, INC.
V.
OMILIA NATURAL LANGUAGE SOLUTIONS, LTD.

CLAIM CONSTRUCTION HEARING

Presentation by Plaintiff Nuance
July 10, 2020

NUANCE

NUANCE

Disputed Term	Nuance's Proposed Construction	Omilia's Proposed Construction
A second language. '925 patent claims 12, 25, 27	A language other than the first language of the first domain/speech recognizer.	A second language not incorporated in the first recognizer.

- **Key Dispute: Whether the "second language" must be absent from the first domain / first speech recognizer**

NUANCE

1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, …
   3. The method of claim 1, …
      4. The method of claim 3…
         6. The method of claim 4, ….
            12. The method of claim 6, wherein said first speech recognizer is a speech recognizer **of at least** a first language and said domain specific training data **relates to** a second language and said second speech recognizer is a multi-lingual speech recognizer **of said second language and said at least first language.**

14. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, …
    16. The machine-readable storage of claim 14, …
        17. The machine-readable storage of claim 16, …
            19. The machine-readable storage of claim 17, …
                25. The machine-readable storage of claim 19, wherein said first speech recognizer is a speech recognizer **of at least** a first language and said domain specific training data **relates to** a second language and said second speech recognizer is a multi-lingual speech recognizer **of said second language and said at least first language.**

27. A computerized method of generating a second speech recognizer comprising the steps of:
- identifying a first speech recognizer of a first domain comprising a first acoustic model with a first decision network and corresponding first phonetic contexts;
- receiving domain-specific training data of a second domain; and
- based on the first speech recognizer and the domain-specific training data, generating a second acoustic model of said first domain and said second domain comprising a second acoustic model with a second decision network and corresponding second phonetic contexts, wherein the first domain **comprises at least** a first language, wherein the second domain **comprises at least** a second language, and wherein the second speech recognizer is a multi-lingual speech recognizer.

"SECOND LANGUAGE" - CLAIM LANGUAGE IS DISPOSITIVE

- Claims 12, 25: First speech recognizer is of "at least" a first language
  - *"At least" means other unclaimed elements may be present*
  - *No language to preclude second language from being present*
- Claims 12, 25: Domain specific training data "relates to" a second language
  - *No requirement that it "add" a second language or that it "is" a second language*
- Claim 27: First domain "comprises" "at least" a first language
  - *"Comprises" is open-ended: includes but is not limited to*
  - *"At least" means other unclaimed elements may be present*
  - *No language to preclude second language from being present*
- Claim 27: Second domain "comprises" "at least" a second language
  - *"Comprises" is open-ended: includes but is not limited to*
  - *"At least" means other unclaimed elements may be present*

NUANCE

CLAIMS 12, 25: FIRST RECOGNIZER MAY INCLUDE SECOND LANGUAGE

"...first speech recognizer is a speech recognizer of at least a first language..."

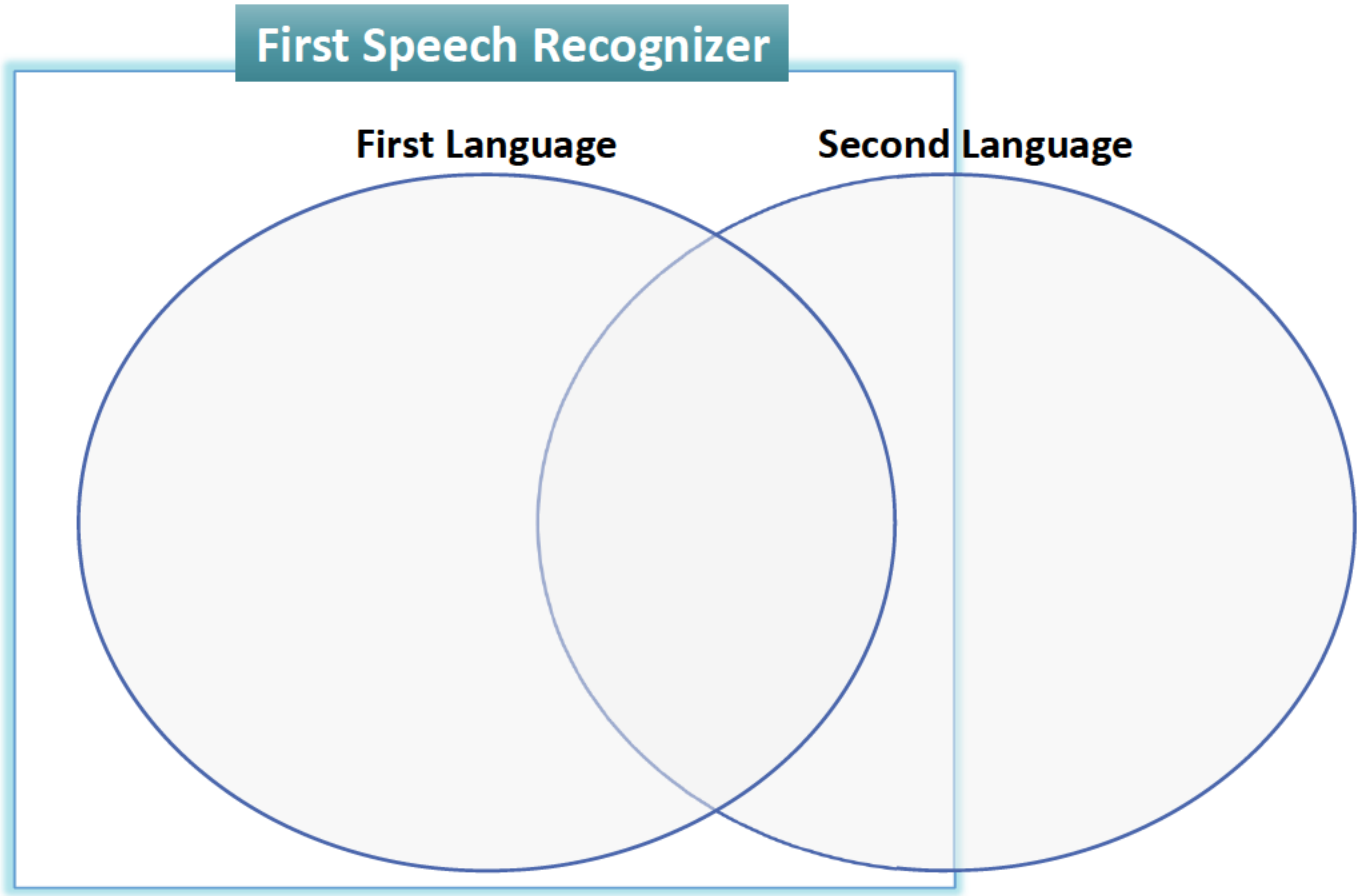

NUANCE

CLAIMS 12, 25: DOMAIN SPECIFIC TRAINING DATA "RELATES TO" (NOT "ADDS") SECOND LANGUAGE

"...and said domain specific training data relates to a second language..."

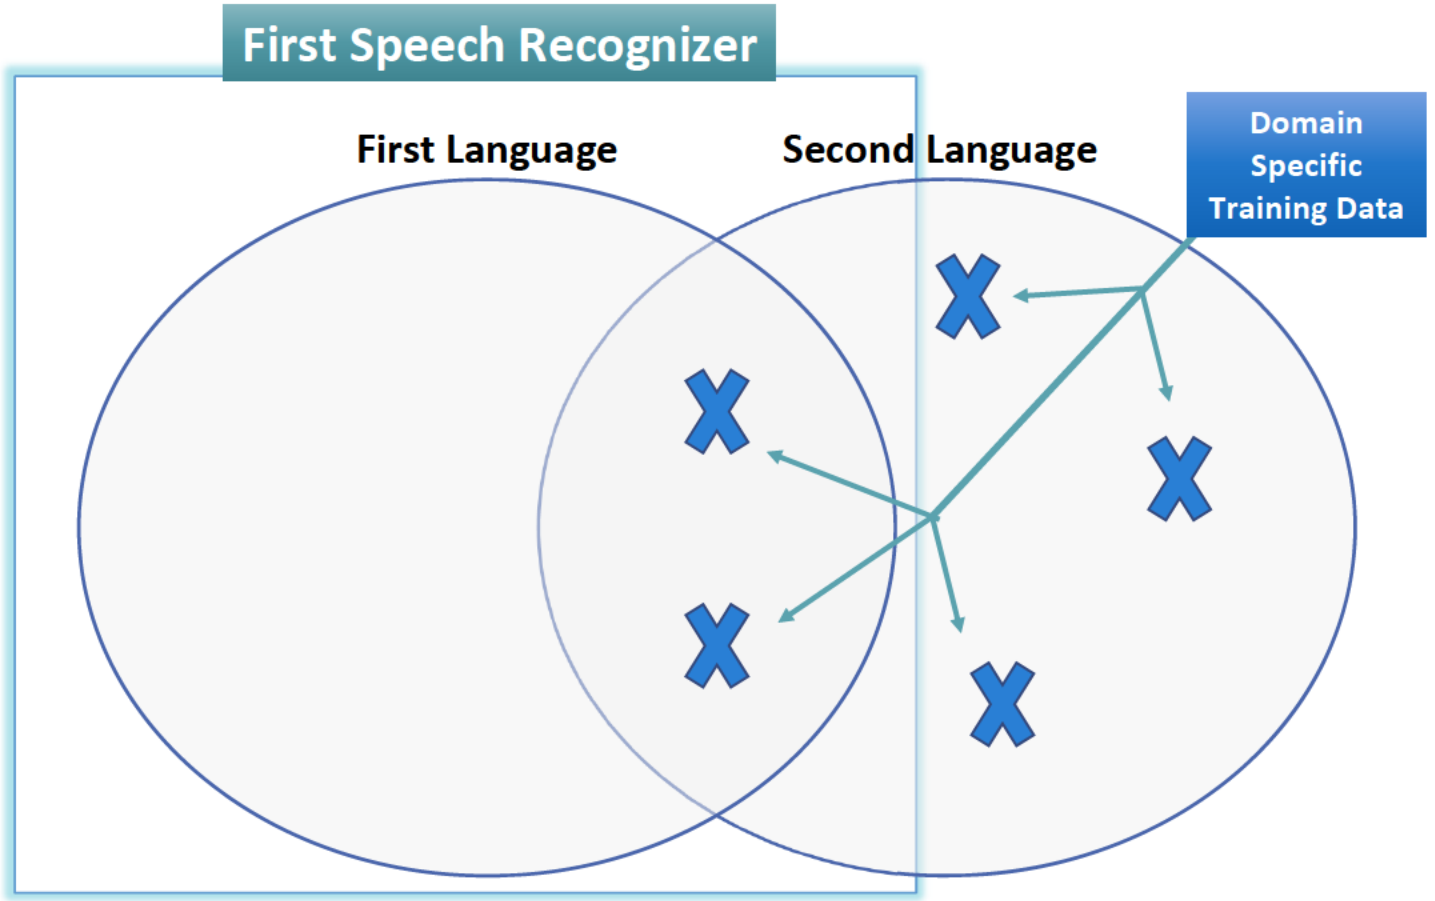

NUANCE

CLAIMS 12, 25: SECOND RECOGNIZER IS MULTI-LINGUAL

"...and said second speech recognizer is a multi-lingual speech recognizer of said second language and said at least first language ..."

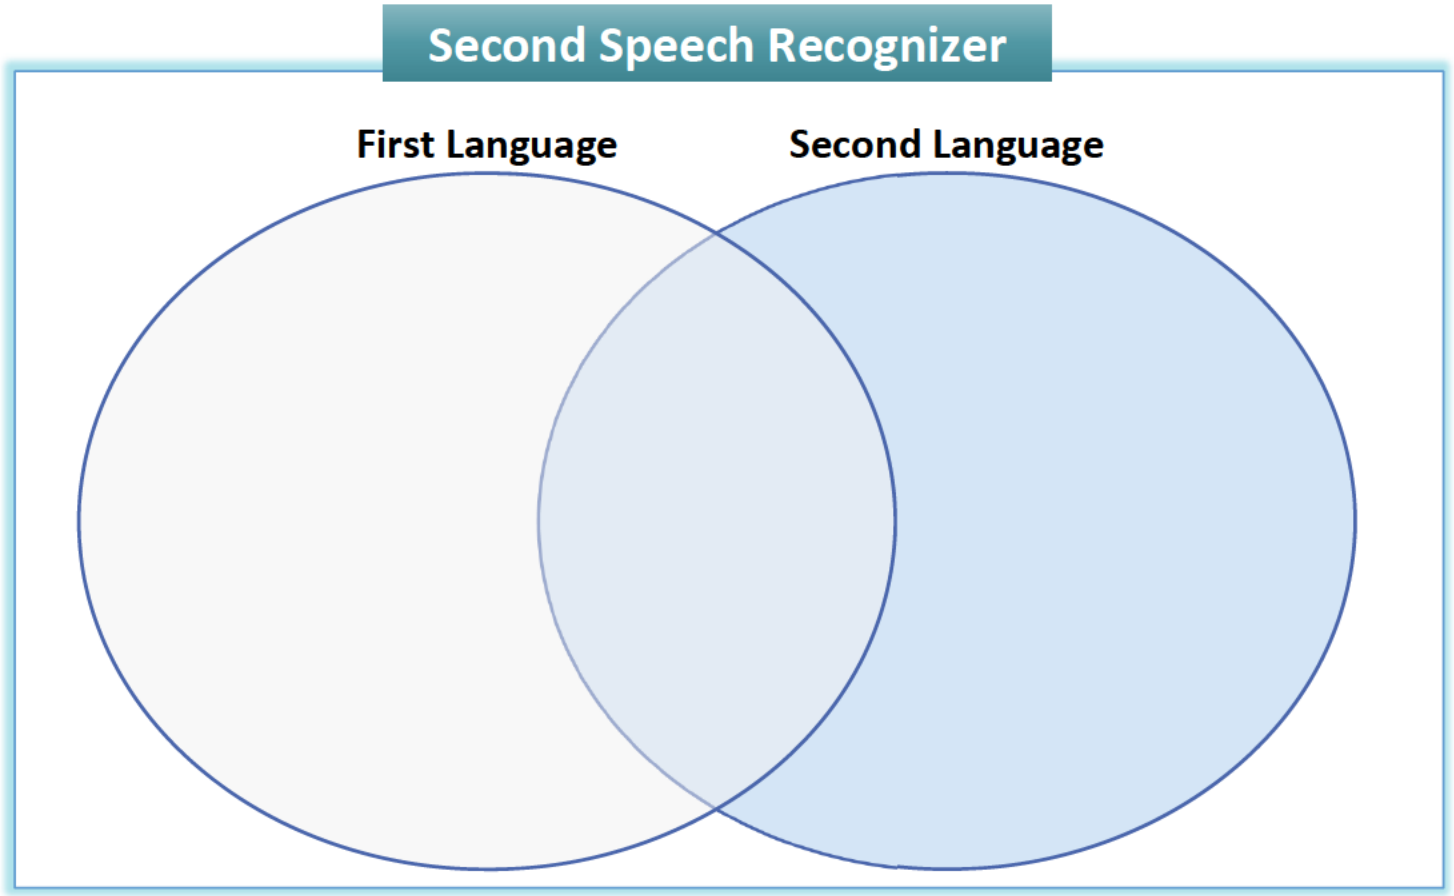

NUANCE

SPECIFICATION, FILE HISTORY, EXTRINSIC EVIDENCE

- Specification
  - Omilia confuses “new domain” with “new language”
  - Domain defined more broadly: languages, dialects, tasks, etc. (col. 6:5-8)
- File History
  - Omilia misreads 11/3/04 Office Action (Sternberg Dec. Ex. 3) and 2/3/05 Response. Waibel rejection is not related to this limitation.
  - Claims 12, 25 (app. claims 13, 27) → allowable as originally drafted if rewritten in independent form
  - Claim 27 (new app. claims 29+30) → allowed as drafted and combined
- Extrinsic Evidence
  - Livescu Decl. ¶¶ 35-40 (languages are not completely distinct, and the claimed invention leverages this for efficient adaptation)

NUANCE

NUANCE

Disputed Term	Nuance's Proposed Construction	Omilia's Proposed Construction
Multi-lingual speech recognizer. '925 patent claims 12, 25, 27	A speech recognizer with an acoustic model that covers sound units of multiple languages.	A speech recognizer incorporating at least the first language and adding the second language.

- **Key Disputes:**
  - **Meaning of "multi-lingual" to a POSITA**
  - **Whether second language must be "added"**

NUANCE

1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, …
3. The method of claim 1, …
4. The method of claim 3…
6. The method of claim 4, ….
12. The method of claim 6, wherein said first speech recognizer is a speech recognizer **of at least** a first language and said domain specific training data **relates to** a second language and said second speech recognizer is a multi-lingual speech recognizer **of said second language and said at least first language.**

14. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, …
16. The machine-readable storage of claim 14, …
17. The machine-readable storage of claim 16, …
19. The machine-readable storage of claim 17, …
25. The machine-readable storage of claim 19, wherein said first speech recognizer is a speech recognizer **of at least** a first language and said domain specific training data **relates to** a second language and said second speech recognizer is a multi-lingual speech recognizer **of said second language and said at least first language.**

27. A computerized method of generating a second speech recognizer comprising the steps of:
- identifying a first speech recognizer of a first domain comprising a first acoustic model with a first decision network and corresponding first phonetic contexts;
- receiving domain-specific training data of a second domain; and
- based on the first speech recognizer and the domain-specific training data, generating a second acoustic model of said first domain and said second domain comprising a second acoustic model with a second decision network and corresponding second phonetic contexts, wherein the first domain **comprises at least** a first language, wherein the second domain **comprises at least** a second language, and wherein the second speech recognizer is a multi-lingual speech recognizer.

EVIDENTIARY SUPPORT

- Claim Language
  - Claims 12, 25: multi-lingual speech recognizer must be "of said second language and said at least first language"
  - Claim 27: domains used to generate multi-lingual speech recognizer must "comprise" "at least" a first language and a second language
- Specification
  - Col. 9:4-23: "If different languages share a common phonetic alphabet, the method also can be used for the incremental and data driven incorporation of a new language into a true multi-lingual speech recognizer …"
- Extrinsic Evidence
  - See Schultz & Waibel (Livescu Dec. Exh. G)
  - POSITA understanding of meaning of "multi-lingual" and presence of a language

NUANCE

PRESENCE OF A LANGUAGE IN A MULTI-LINGUAL SPEECH RECOGNIZER

- Each language has a number of phonemes
- Ability to recognize the phonemes means the language is supported
- More training data / specialized training data → more accurate recognition
- "Sound units" in Nuance's proposed construction, as understood by a POSITA, includes acoustic contexts (Nuance Resp. Br. at 3-4)

- Omilia's misconceptions
  - A language is either present or absent
  - Language is an all or nothing exercise
  - A phonetic context is the same as an entire language

- See examples of multi-lingual speech recognizers in Livescu Ex. G

NUANCE

SCHULTZ & WAIBEL (LIVESCU DECL. EXH. G)

Phonemes [Worldbet]	KO	SP	CR	TU	JA	∑
n,m,s,l,tS,p,b,t,d,g,k	X	X	X	X	X	
i,e,o	X	X	X	X	X	14
f,j,z		X	X	X	X	
r,u	X	X	X	X		
dZ	X		X	X	X	6
a	X	X	X			
S			X	X	X	
h	X			X	X	
4	X	X			X	4
$\tilde{n}$,x,L		X	X			
A				X	X	
N	X	X				
V,Z			X	X		
y,7	X			X		
ts			X		X	10
p',t',k',dZ',s',oE,oa,4i,	X					
uE,E,Λ,iΛ,uΛ,iu,ie,io,ia	X					17
D,G,T,V,r(,ai,au,ei,eu,oi		X				
a+,e+,i+,o+,u+		X				15
palatal c, palatal d			X			2
ix, soft				X		2
?,Nq,V[,A:,e:,i:,o:,4:					X	8
Monolingual $\sum = 170$	40	40	30	29	31	
Multilingual						78

Table 3: Global Phoneme Set [Worldbet notation]

NUANCE

THREE MULTILINGUAL SPEECH RECOGNIZER SYSTEMS

Color Key

- Korean (KO)
- Spanish (SP)
- Croatian (CR)
- Turkish (TU)
- Japanese (JA)

Phonemes [Worldbet]	KO	SP	CR	TU	JA	∑
n,m,s,l,tS,p,b,t,d,g,k	X	X	X	X	X	
i,e,o	X	X	X	X	X	14
f,j,z		X	X	X	X	
r,u	X	X	X	X		
dZ	X		X	X	X	6
a	X	X	X			
S			X	X	X	
h	X			X	X	
4	X	X			X	4
ñ,x,L		X	X			
A				X	X	
N	X	X				
V,Z			X	X		
y,7	X			X		
ts			X		X	10
p',t',k',dZ',s',oE,oa,4i,	X					
uE,E,ʌ,iʌ,uʌ,iu,ie,io,ia	X					17
D,G,T,V,r(,ai,au,ei,eu,oi		X				
a+,e+,i+,o+,u+		X				15
palatal c, palatal d			X			2
ix, soft				X		2
?,Nq,V[,A:,e:,i:,o:,4:					X	8
Monolingual $\sum = 170$	40	40	30	29	31	
Multilingual						78

Table 3: Global Phoneme Set [Worldbet notation]

NUANCE

** Exhibit G of Livescu Decl.*

MULTILINGUAL SYSTEM *ML5-MIX*

78 PHONEMES, LANGUAGES MIXED

n,m,s,l,tS,p,b,t,d,g,k

i,e,o

a

S

h

4

$\tilde{n}$,x,L

A

N

V,Z

y,7

ts

p',t',k',dZ',s',oE,oa,4i,

uE,E,Λ,iΛ,uΛ,iu,ie,io,ia

D,G,T,V,r(,ai,au,ei,eu,oi

a+,e+,i+,o+,u+

palatal c, palatal d

ix, soft

?,Nq,V[,A:,e:,i:,o:,4:

f,j,z

r,u

dZ

NUANCE

** Exhibit G of Livescu Decl.*

MULTILINGUAL SYSTEM *ML5-SEP*

EACH PHONEME MODELED SEPARATELY FOR EACH LANGUAGE

Color Key

- Korean (KO)
- Spanish (SP)
- Croatian (CR)
- Turkish (TU)
- Japanese (JA)

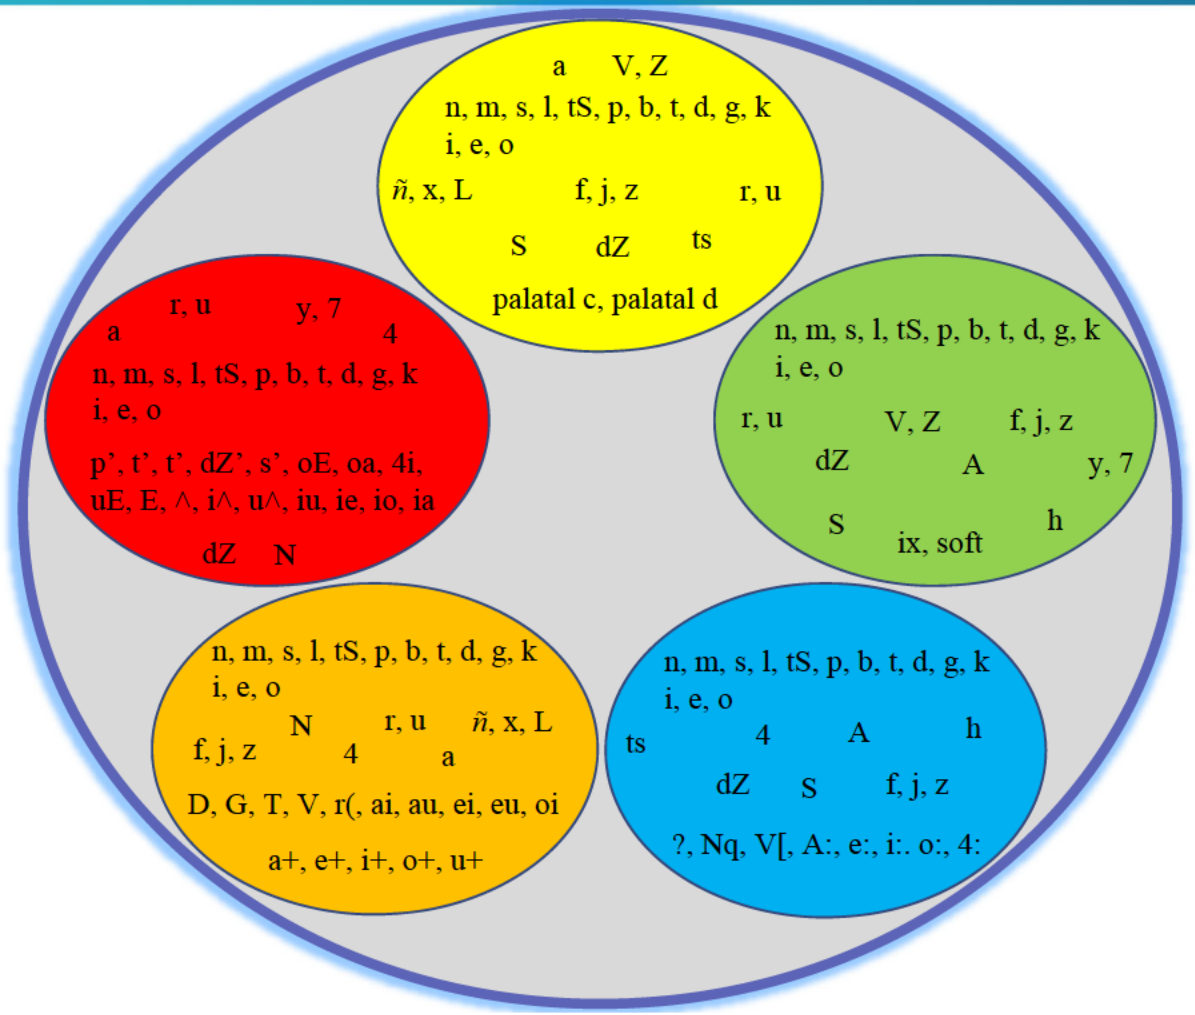

NUANCE

** Exhibit G of Livescu Decl.*

MULTILINGUAL SYSTEM *ML5-TAG*

EACH PHONEME TAGGED WITH LANGUAGE IDENTIFIER

Color Key
Korean (KO)
Spanish (SP)
Croatian (CR)
Turkish (TU)
Japanese (JA)

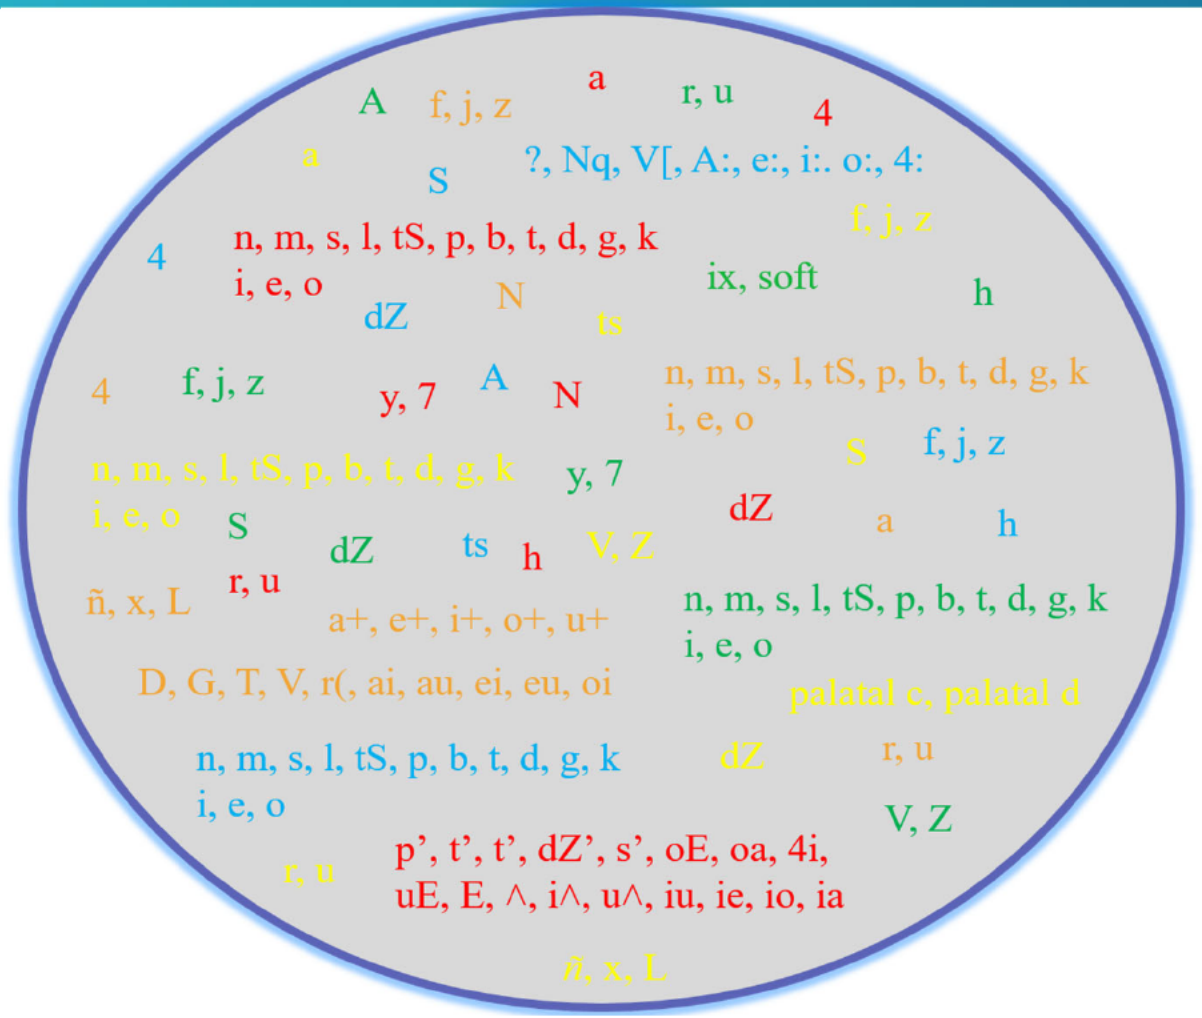

NUANCE

** Exhibit G of Livescu Decl.*

NUANCE

Disputed Term	Nuance's Proposed Construction	Omilia's Proposed Construction
Generating a second acoustic model. '925 patent claim 27	Plain and ordinary meaning.	"generating a second acoustic model" requires "re-estimating said first decision network and said corresponding first phonetic contexts"

- **Key Dispute: Whether "generating" requires "re-estimating said first decision network and said corresponding first phonetic contexts"**

NUANCE

1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision free to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises partitioning said training data using said first decision network of said first speech recognizer.	14. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises partitioning said training data using said first decision network of said first speech recognizer.	27. A computerized method of generating a second speech recognizer comprising the steps of: identifying a first speech recognizer of a first domain comprising a first acoustic model with a first decision network and corresponding first phonetic contexts; receiving domain-specific training data of a second domain; and based on the first speech recognizer and the domain-specific training data, generating a second acoustic model of said first domain and said second domain comprising a second acoustic model with a second decision network and corresponding second phonetic contexts, wherein the first domain comprises at least a first language, wherein the second domain comprises at least a second language, and wherein the second speech recognizer is a multi-lingual speech recognizer.

NUANCE

2. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model wit a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of: . . .	15. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of: . . .	27. A computerized method of generating a second speech recognizer comprising the steps of: identifying a first speech recognizer of a first domain comprising a first acoustic model with a first decision network and corresponding first phonetic contexts; receiving domain-specific training data of a second domain; and based on the first speech recognizer and the domain-specific training data, generating a second acoustic model of said first domain and said second domain comprising a second acoustic model with a second decision network and corresponding second phonetic contexts, wherein the first domain comprises at least a first language, wherein the second domain comprises at least a second language, and wherein the second speech recognizer is a multi-lingual speech recognizer.

"GENERATING" – CLAIM LANGUAGE IS DISPOSITIVE

- Claims 1, 2, 14, 15 all require "generating … by re-estimating …"
  - Omilia's proposed construction would render the "by re-estimating" limitation redundant
- Claim 27 requires "generating" but does not include "by re-estimating"
  - Omilia's proposed construction would import limitation language into claim 27 that is not present

NUANCE

SPECIFICATION, FILE HISTORY, EXTRINSIC EVIDENCE

- Specification
  - Acknowledges other approaches possible. *See* Nuance Br. at 15 (cites)
  - Nothing in specification is *definitional* as to "generating"
- File History
  - Claim 27 allowed without requiring "by re-estimating …" (Livescu Decl. Exhs. J, I, J)
- Extrinsic Evidence
  - Livescu Dec. ¶¶ 50-53 (how a POSITA would understand claims, specification and file history)

NUANCE

NUANCE

Disputed Term	Nuance's Proposed Construction	Omilia's Proposed Construction
Automatically generate/ing. '925 patent claims 1, 2, 14, 15	Generate/ing, at least in part by a computer.	Generate/ing by a computer without human intervention.

- **Key Issue: Whether "automatically generating" precludes any human participation as to unclaimed steps**

NUANCE

1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision free to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises partitioning said training data using said first decision network of said first speech recognizer.	14. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, and wherein said re-estimating comprises partitioning said training data using said first decision network of said first speech recognizer.

NUANCE

2. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model wit a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of: . . .	15. A machine-readable storage, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said machine-readable storage causing the machine to perform the steps of: based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data, wherein said first decision network and said second decision network utilize a phonetic decision tree to perform speech recognition operations, wherein the number of nodes in the second decision network is not fixed by the number of nodes in the first decision network, wherein said domain-specific training data is of a limited amount, and wherein the generating step further comprises the steps of: . . .

DISPUTE IS OVER HUMAN INVOLVEMENT IN UNCLAIMED STEPS

- Nuance agrees that the claim requires that the claimed steps are performed by a computer.
  - "Because the 'automatically generating' term appears in the preamble, it requires that each *claimed step* of the method must be automatically performed." Omilia Resp. Br. at 13.
- Other, unclaimed, actions or steps may be involved. The claim is silent on whether those unclaimed steps must be performed by a computer.
  - *E.g.*, collecting data, setting thresholds or parameters. Livescu Decl. ¶¶ 54-59.
  - *See* Fig. 1 (supervised collection of application-specific training speech data)
- Omilia's Responsive Brief at 13
  - "Omilia's construction does not preclude that other steps may be performed by humans."
  - Omilia's proposed construction "does not mean that a human may never be involved."

NUANCE

925 PATENT INCLUDES BOTH METHOD AND COMPUTER READABLE MEDIUM CLAIMS

- Claims 14-15
  - "**A machine-readable storage**, having stored thereon a computer program having a plurality of code sections executable by a machine for causing the machine to automatically generate … a second speech recognizer, … said machine-readable storage **causing the machine to perform the steps of**: …"
    - *Omilia's construction would require the software to be programmed to perform other, unclaimed, steps as well. This improperly imports limitations into the claims.*
- Claims 1-2
  - "**A computerized method** of automatically generating … a second speech recognizer, … said method **comprising**: …
    - *To require that "automatically generating" in claims 1-2 include computer-performance of steps other than those specified would improperly (a) import limitations into the claim, and/or (b) apply a different meaning for "automatically generating" to different claims.*

NUANCE